Date: Monday, 21 July 2014
Facility: Fairmont Royal York Hotel, Ontario, Canada
Room: Manitoba Room
Time: 09:00 - 11:30

9:05 ó 1:30
5 min. Administrivia (scribes etc.), Agenda Bashing, Chairs

Please note the note well slide. If you know IPR, you should report it. Sue Hares will be scribing. Donald read the Note-Well.
Donald asked for a Jabber scribe but there was no volunteer.

8 min. Document Status,
    Review of Milestones, Chairs
        (see the slides)

Draft-ietf-trill-rbridge-vlan-mapping will be allowed to expire unless there is some objection. There seems to have been little interest in working on this draft for some time.

5 min. TRILL Implementation Report, Susan Hares

9:18:
15 min. TRILL OAM Status Update,
    Tissa Senevirathne, Tal Mizrahi, Deepak Kumar

  draft-ietf-trill-oam-fm-06,
  draft-ietf-trill-loss-delay-05
 draft-ietf-trill-oam-mib-00
  draft-tissa-yang-oam-00

Major Work:
The Co-Chairs believe that the relative priority of the current major specification efforts in the TRILL WG is as below. Major efforts typically have multiple documents.
  o OAM
  o Active-Active
  o ARP & ND optimization / Directory Assist

Discussion [At slide 3]
• Sue: What are you looking for in Yang?
• Tissa: We∅ll cover this in a moment.
[Ted Lemon: Sorry to delay on a process issue, but we really shouldn∅t be allowing folks to ask questions without the mic as it won∅t get recorded.

(Mic setup changed)]

- Weigou: Your generic model is for all overlay networks.
- Tissa: I am not targeting any specific overlays. From our point of view overlays and underlays donøt really matter, we want to provide the framework to provide any architecture so that any overlay like TRILL or NVo3 can easily map to it. The general model presents an overlay framework that will bound the OAM model. The only OAM model which exists is the IEEE CFM model. This document creates a framework or architecture. The model underneath will support. For example, you will provide the TRILL specific underneath ó it will be
- Weigo: It is a large architecture.
- Tissa: The flexibility will provide the ability. The future OAM developer architecture document will be discussed in the Operations meeting. The unified management framework will from Yang will provide a unified mange
- Sue: Will the Yang modules at the top have the templates (yang groups) have the general definition?
- Tissa: This is the design for the OAM. The TRILL will pull in the OAM Yang modules for the TRILL. I will show more detail in the presentation.
- Sue Hares: Does your Yang model offer generic templates so that for example TRILL could use it without significant work?
- Tissa: Yes this is the goal
- Sue: This works well if the OAM is the top node. How does this work if the TRILL is the top node for configuration, i2rs, and OAM.
- Tissa: We need a top-level.

Discussion at the end of the presentation:

- Tissa: Next Steps? Should we adopt draft-tissa-yang-oam-00
- Frank Xia(Huawei): Why is there a separate Yang mode for Loss and Delay for TRILL? Can you explain this? Why will you do this?
- Tissa: Keep it simple to make Loss and Delay for TRILL. Some people will not implement the yang model for Loss and Delay for TRILL. The Yang model will become? IT can be combined or separate?
- Frank Xia : My concern is that have the base document without the loss and delay would miss some functions.
- Tissa: All the functions exist in the TRILL model. Some of Loss and delay. Yang allows a sub-module.
- Sue: I think this should be a sub-module that can be attached to the TRILL OAM model. You should
- Trissa: We are in violent agreement.
- Weigo: You want to introduce a generic model for Loss and delay?
- Tissa: Trill has both loss and delay and OAM. Most areas do not have loss and delay.
- Weigo: You think Loss and delay is a specific area issue, it shouldnøt be included in generic Yang model.
- Tissa: Loss and delay is an operation under generic OAM rather than a top-level model. The top level OAM can be extended to NVO3.

- Donald: Co-chairs will post a adoption call today for the yang-oam draft running for three weeks as part of it overlaps with this IETF meeting.

35 min. TRILL Active-Active,
      Hongjun Zhai, Tissa Senvirathne, Yizhou Li,
      Weiguo Hao, Mingui Zhang

        draft-hu-trill-pseudonode-nickname-08
    draft-zhang-trill-aa-multi-attach-04
    draft-ietf-trill-cmt-03
    draft-hao-trill-analysis-active-active-01

Donald Eastlake presented pseudonode-nickname for Yizhou who missed an airplane connection and won¢t be at the IETF meeting until tomorrow.   Mingui Zhang presented aa-multi-attach.

Discussion:

- Andrew Qu (Media-Tek): Hardware will have to change to be able to rewrite the local nickname to the virtual nickname of the MC-LAG. I know of no hardware that can support this. Second, I think the there may be confusion in the terminology

- Donald: I am not familiar with all the hardware.

- Andrew: We are working on hardware. The source nickname will have functionality that will change the packet, and then it comes around a second time. I¢m not sure if the source nick name is the right term.

- Donald: the ingress nickname is the right way. You need to check if the
- Andrew: It is confusing to call it a source bridge.

- Donald: Source bridge and nick name denote the same thing.

- Andrew: I do not think hardware can do this function at this point.

Discussion:
- Weiguo: The member Rbridge should always be discovered. I think this should be separated from the draft.
- Donald: The mechanisms in the two drafts are the same. We should be able to factor these mechanisms out into a common draft.
- Weiguo: How can we avoid the traffic loop among this solution?
- Mingui: In RB3, there will be 2 roots.
- Weiguo: Do you want to use single nickname for split-horizon?
- Mingui: the operators use their own nickname, not a pseudo-nickname.

- Weiguo: OK, you don't rely on pseudo-nickname. Local bias forwarding behavior should be used. If CE1 and CE2 devices are both directly attached to RB1 and RB2, when RB1 receives the traffic from CE1, RB1 should forward the traffic to CE2 locally no matter local port DF state. When RB2 receives the traffic from TRILL network side, it drops the traffic relying on ingress nickname.
- Donald: This technical discussion can be best described on the list.
- Andrew: I saw earlier you had unique RB4 have active attachment to RB4, but here you reduced to have one RBridge or the other. Is there anyway for the ingress to remember them all?
- Donald: You remember them all, and only install one in the FIB.
- Andrew: How do you know distribution on the traffic is balanced?
- Donald: It is the law of large numbers. Probabilistically it will spread more or less evenly.
- Andrew: What happens if the packets are sent?
- Donald: If they are pseudo random (as in the ESADI draft now in the RFC Editor queue).
- Andrew: What if the packets are not randomly distributed?
- Donald: Look at section 5.4 of the ESADI draft.

Donald:
- Based on comments here and on mail list, we will adopt both of these drafts as WG drafts and update the terminology to be as in draft-ietf-trill-active-active-connection-prob.

12 min. TRILL over IP, Margaret Wasserman, Dacheng Zhang
  draft-ietf-trill-over-ip-01

Discussion:
- Ted Lemon: How are you handling congestion control?
- Donald: The congestion provisions in the current draft were stolen from an old version of the MPLS over UDP IP draft. Experience indicates that people will run over the UDP and it just works. There was a comment on the list to pay attention to the ___ RFC for congestion control. The problem occurs if the link are shared. You could plug both pieces of the TRILL infrastructure and it will just work. If you exceed the traffic limits per link.
- Andrew: For TRILL over IP, our product already uses VXLAN. We added it for common deployment scenarios. The VXLAN carries between two TRILL subnets. Why not use the VXLAN?
- Donald Qu: what is the VxLAN running over?
- Andrew: The TRILL can use UDP (VxLAN)over ethernet. I can then differentiate between the native TRILL and the data packets. The VXLAN can be very flexible.
- Donald: there are multiple ways of connecting between parts of a campus. I claim there should be one standard way to connect between the two campuses for interoperability. I just suggest that UDP is a simply way to run over IP.
- Andrew: the VxLAN has UDP packet frames.
- Donald: It is just another Ethernet header after the vXLAN header.

- Andrew Qu: When you decapsulate from the packet, you need another decapsulation? Should we consider using VXLAN instead of IP?
- Donald: We have draft-ietf-trill-over-ip-01.txt
- Ted: The congestion provisions in the current draft do not look adequate.
- Donald: Yes they are deficient but there is a new version of the MPLS over UDP draft with improved text we could probably use.
- Ted: OK, Thanks.

12 min. Directory Assisted Edge Update,
 Donald Eastlake, Linda Dunbar
draft-ietf-trill-directory-assist-mechanisms-00
draft-ietf-trill-ia-appsubtlv-00
draft-ietf-trill-chanel-tunnel-00
draft-dunbar-trill-directory-assisted-encap-07

Discussion:
- Does anyone have problem splitting section 5 and parts of Section 4 out from the draft-ietf-trill-directory-assistance-mechanisms?
- How many people have read the draft?
  [4 or 5 people raised their hands].
- Donald: That's not too terrible, but not too good. Does anyone have a problem splitting section 4/5 away from the assist-mechanisms draft?
- no comments.
- Donald: We'll take it to the last.

22 min. Layer 3 Gateway: Weiguo Hao
Distributed: draft-aho-trill-irb-04
Anycast: draft-hao-trill-anycast-gw-00

Discussion on draft-hao-trill-irb:
- Weiguo: I would alike to ask for WG adoption.
- Donald: We should take this to the list. We should wait a few weeks and then do WG adoption.

Discussion on anycast-gw:
- Andrew: Anycast is put on every port. You did not cover this issue.   Each port handling the
- Each TOR can use the anycast to the TOR. I can choose to use the
- aggregation to the TOR. If you use two gateway, then this would
- If you anycast for the same IP host entry, it will be best.
- It will be important to the know the best local lan account

Weiguo:

Weiguo: I think the distributed gateway is for east-west direction traffic forwarding gateway. Anycast gateway is centralized gateway, it is used for north-south direction traffic forwarding.

Andrew: for example, for the VM-move the gateway is not chosen.

Weiguo: You should remove the local gateway for VM-move case.

Andrew: VM-move are trying to send the Virtual-Mac/Virtual-IP. The Virtual IP address should be moved to all TORs. You need the IP addresses distributed to all Gateways. You should distribute the gateway to all ports. You need the x port distributed to the port. You need to cover the local gateway. We can discuss this off line.

12 min. Further TRILL Clarifications, Corrections, and Updates
rfc7180bis, Donald Eastlake

Discussion:
Ted: What did you say about this being a personal draft?
Donald: I just said that the plan is for rfc7180bis to start as a personal draft and go through the usual process of adoption.
Ted: OK.

4 min. Wrap-Up, Chairs

Donald: Thank you for participating. See you in Honolulu.

11:26 adjourn.